Introducing IBM Db2 for z/OS Data Gate

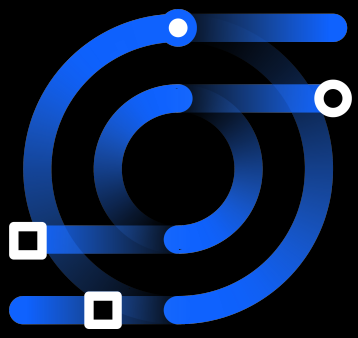

Executive Summary

Access Db2 to z/OS data without accessing Db2 for z/OS

IBM Db2 for z/OS Data Gate (Db2 Data Gate) offers a new way for IBM Z customers to deliver data originating in Db2 for z/OS. It enables organizations to access Db2 for z/OS data without accessing Db2 for z/OS—and without impact to the source systems on IBM Z. Db2 Data Gate provides an integrated approach to synchronize and deliver data originating in Db2 for z/OS to support the increasing demand for analytics and dramatic increase in modern high-volume, high frequency read-only applications.

Hybrid Cloud deployment and data access

Db2 Data Gate leverages IBM Cloud Pak for Data, a fully integrated data and AI platform that modernizes how businesses collect, organize and analyze data to infuse data-powered AI throughout their organizations. Data is synchronized between Db2 for z/OS data sources and target databases on IBM Cloud Pak for Data.

IBM Cloud Pak for Data is optimized to integrate within hybrid cloud deployment models including on-premises, private and public clouds, and in pre-integrated systems. Db2 Data Gate therefore, facilitates the use of data originating in Db2 for z/OS within all hybrid cloud scenarios provided by Cloud Pak for Data, enabling a hybrid cloud consumption model for Z data in an efficient, integrated and easy to use way.

Introducing IBM Db2 for z/OS Data Gate

New and modernized applications

Digital transformation is offering unprecedented opportunities for organizations to improve their customer relationships, acquire new customers and grow wallet share. Organizations are innovating with new ways to engage customers and take advantage of their enterprise data assets. Data from traditional applications have become the lifeblood of new cloud, analytics and mobile applications. These new applications are driving significantly greater IT activity.

Organizations want to leverage the valuable data that originates on IBM Z for information-driven projects that run on non-IBM Z platforms. Many organizations have embarked on projects that attempt to copy data from Db2 for z/OS to various targets in real-time. These targets include traditional relational database management systems, data warehouses, Hadoop data lakes, Kafka streaming hubs, and others.

Modernization efforts are often associated with making new, more engaging, flexible applications available to an information consumer. These new engagement applications simplify tasks for the customer and often minimize the need for customer service intervention with everyday tasks. New engagement applications are aligned to a generation that is comfortable with internet and mobile devices and would rather do things themselves than interact with a customer service agent. Modern applications allow organizations to reach new demographics and expand access to new customers less expensively.

IBM

This surge of new digital applications is driving an exponential growth in data access. These new applications often leverage read-only lookup type queries of up-to-date data. Alternatively, they can require complex queries executed against large sets of data. These applications can be seasonal and unpredictable with significant spikes in usage, making them difficult to plan for and size appropriately. Many IT organizations consider these applications necessary to support the business but would understandably like to minimize their associated cost. Yet, as these applications are often internal and external customer facing, it is vitally important that the data be current and access to that data needs to be resilient.

The volume of queries from these new engagement applications continues to increase. The queries often originate via mobile, cloud and self–service applications. They are also frequently accompanied by updates to logging or activity-tracking tables. Changes to source transactional data only occurs at the source system databases. Addressing these application needs with home-grown extract and load approaches can be associated with a high degree of data latency. These in-house approaches can limit efficiency and performance improvement opportunities. Data extraction and incremental copy processes can consume considerable resources, drive up related software costs and compete for the same resources used for operational processing.

These new customer facing applications at times require extensive data management efforts to maintain data currency. As these applications are customer facing often the data must be up to date, only moments behind a transactional system. This has driven the need for more complex application logic ensuring data currency. On top of the processes to refresh the read-only data store, many organizations have used customized code to ensure consistency. This has added more complexity, instability and cost to many environments.

Db2 Data Gate Overview

Db2 Data Gate provides an integrated, resilient, standardized approach, designed to be more reliable, less complex and less expensive than customized solutions. It supports applications that require read-only access to Db2 for z/OS data.

Data access patterns can be similar to transactional access paths such as relatively simple index-addressable read-only point queries. They can also be similar to data warehousing queries requiring complex data-intensive query processing. A single Db2 Data Gate repository can be optimized for either high performance index-addressable read-only point queries or high-performance complex data-intensive query processing.

Multiple instances of Db2 Data Gate repositories can be used to respond to both access patterns simultaneously.

Db2 Data Gate leverages IBM Cloud Pak for Data to deliver a synchronized copy of the original Db2 for z/OS data. An application accesses a copy of the Db2 for z/OS data instead of accessing data directly in Db2 for z/OS. The Db2 Data Gate engine manages data on a separate Linux-based system as part of IBM Cloud Pak for Data.

Access to Db2 Data Gate will be offered directly through modern SQL APIs. The request does not get passed through Db2 for z/OS.

Db2 for z/OS is the core transactional database engine used by IBM Z clients throughout the world. With Db2 Data Gate, Db2 for z/OS remains the core transactional processing engine storing and consistently managing critical data for core applications.

Db2 Data Gate synchronizes data from Db2 for z/OS to make it securely available within IBM Db2 Warehouse on IBM Cloud Pak or IBM Cloud Pak for Data Db2. Db2 Warehouse supports data-intensive, complex query processing while Cloud Pak for Data Db2 supports high-frequency, simple read requests. Clients have a choice of targets for Db2 for z/OS data based on specific application requirements.

Db2 Data Gate provides an end-to-end solution to ensure that Db2 for z/OS data is available and synchronized from sources on IBM Z to targets optimized on IBM Cloud Pak for Data. A log data provider captures log changes from Db2 for z/OS and sends consolidated, encrypted changes to a log data processor residing on IBM Cloud Pak for Data. These changes can be sent to one or more Db2 Data Gate targets, reducing the complexity and cost of application development. The log capture processing is fully zIIP enabled to ensure that there is little impact to general processing on IBM Z. Targets are determined based on consuming application requirements and are fully optimized to ensure low latency and high throughput.

An integrated, patented synchronization protocol maintains data currency between the source system on IBM Z and Cloud Pak for Data Db2 or Db2 Warehouse on IBM Cloud Pak for Data. This low-latency, high-throughput synchronization protocol contains numerous optimization features. Data in the target systems on IBM Cloud Pak for Data can be updated and new data inserted, but none of these changes are propagated back to Db2 for z/OS, which retains data ownership.

Summary

IBM Db2 for z/OS Data Gate V1.1 delivers synchronized data from Db2 for z/OS to targets on IBM Cloud Pak for Data with the following key capabilities:

- Db2 Data Gate provisions and enables a service on IBM Cloud Pak for Data
- A fully zIIP enabled, integrated synchronization feature maintains currency between source Db2 for z/OS data on IBM Z and targets on Cloud Pak for Data
- Target databases can be Db2 Warehouse on IBM Cloud Pak for Data or Cloud Pak for Data Db2, relational databases that deliver advanced data management and analytics capabilities

Leveraging IBM Db2 for z/OS Data Gate, you can drive greater value from your existing IBM Z investment. It allows you to keep pace with new cloud, analytics and mobile initiatives with reduced cost and effort. It can help you reduce the complexity associated with delivering data for these new initiatives. With IBM Db2 for z/OS Data Gate, your infrastructure investments can accelerate your hybrid cloud journey.

© Copyright IBM Corporation 2020. IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.